

Chapter 2

The Linear Integration Problem I

In the first part of the book, we consider the linear integration problem **I** and establish a comparison with its $(\max, +)$ -algebra analogue, the linear programming (LP) problem **P**.

2.1 Introduction

In this chapter, we are interested in the *linear integration* problem **I** defined in Chapter 1, that is,

$$\mathbf{I} : \widehat{f}(y, c) := \int_{\Omega(y)} e^{c'x} d\sigma, \quad (2.1)$$

where $\Omega(y) \subset \mathbb{R}^n$ is the convex polyhedron

$$\Omega(y) := \{x \in \mathbb{R}^n \mid Ax = y, \quad x \geq 0\} \quad (2.2)$$

for some given matrix $A \in \mathbb{R}^{m \times n}$, and vectors $y \in \mathbb{R}^m$, $c \in \mathbb{R}^n$, all with rational entries, and where $d\sigma$ is the Lebesgue measure on the smallest affine variety that contains $\Omega(y)$. When $c \equiv 0$, A is full rank, and $\Omega(y)$ is compact, we thus obtain the $(n - m)$ -dimensional volume of $\Omega(y)$. As already mentioned, the linear program **P** is the $(\max, +)$ -algebra analogue of problem **I** in the usual algebra $(+, \times)$, and we will see in the next chapter that the analogy between **I** and **P** is not simply formal.

We describe what we call a *dual* algorithm for computing the volume of $\Omega(y)$, i.e., problem **I** with $c \equiv 0$, because it is a standard problem in computational geometry with many interesting and important applications, and also because the same algorithm works when $c \neq 0$, with *ad hoc* and obvious modifications. Also it can be a viable alternative or complement to the various primal methods briefly presented below.

Some computational complexity results

Computing the volume $\text{vol}(\Omega)$ of a convex polytope Ω (or integrating over Ω) is difficult; its computational complexity is discussed in, e.g., Bollobás [24] and Dyer and Frieze [46]. Indeed, any deterministic algorithm with polynomial time complexity that would compute upper and lower bounds \bar{v} and \underline{v} on $\text{vol}(\Omega)$ cannot yield an upper bound $g(n)$ on \bar{v}/\underline{v} better than a polynomial in the dimension n .

A *convex body* $K \subset \mathbb{R}^n$ is a compact convex subset with nonempty interior. A strong separation *oracle* answers either $x \in K$ or $x \notin K$, and in the latter case produces a hyperplane separating x from K . An algorithm is a sequence of questions to the oracle, with each question depending on the answers to previous questions. The complexity of the algorithm is the number of questions asked before upper and lower bounds $\bar{\text{vol}}(K)$ and $\underline{\text{vol}}(K)$ are produced. Let $B^n \subset \mathbb{R}^n$ be the Euclidean unit ball of \mathbb{R}^n . If $r_1 B^n \subset K \subset r_2 B^n$ for some positive numbers r_1, r_2 , the algorithm is said to be *well guaranteed*. The input size of a convex body satisfying $r_1 B^n \subset K \subset r_2 B^n$ is $n + \langle r_1 \rangle + \langle r_2 \rangle$, where $\langle x \rangle$ is the number of binary digits of a dyadic rational x .

Theorem 2.1. *For every polynomial time algorithm for computing the volume of a convex body $K \subset \mathbb{R}^n$ given by a well-guaranteed separation oracle, there is a constant $c > 0$ such that*

$$\frac{\bar{\text{vol}}(K)}{\underline{\text{vol}}(K)} \leq \left(\frac{cn}{\log n} \right)^n$$

cannot be guaranteed for $n \geq 2$.

However, Lovász [107] proved that there is a polynomial time algorithm that produces bounds $\bar{\text{vol}}(K)$ and $\underline{\text{vol}}(K)$ satisfying $\bar{\text{vol}}(K)/\underline{\text{vol}}(K) \leq n^n (n+1)^{n/2}$, whereas Elekes [54] proved that for $0 < \varepsilon < 2$ there is no polynomial time algorithm that produces $\bar{\text{vol}}(K)$ and $\underline{\text{vol}}(K)$ with $\bar{\text{vol}}(K)/\underline{\text{vol}}(K) \leq (2 - \varepsilon)^n$.

In contrast with this negative result, and if one accepts randomized algorithms that fail with small probability, then the situation is much better. Indeed, the celebrated Dyer, Frieze, and Kanan's probabilistic approximation algorithm [47] computes the volume to fixed arbitrary relative precision ε , in time polynomial in ε^{-1} . The latter algorithm uses approximation schemes based on rapidly mixing Markov chains and isoperimetric inequalities.

More precisely, following [24], let $K \subset \mathbb{R}^n$ with $n \geq 2$, and let ε, η be small positive numbers. An ε -approximation of $\text{vol}(K)$ is a number $\widehat{\text{vol}}(K)$ such that

$$(1 - \varepsilon) \widehat{\text{vol}}(K) \leq \text{vol}(K) \leq (1 + \varepsilon) \widehat{\text{vol}}(K). \quad (2.3)$$

A *fully polynomial randomized approximation scheme* (FPRAS) is a randomized algorithm that runs in time polynomial in the input size of K , ε^{-1} , $\log \eta^{-1}$, and with probability at least $1 - \eta$ produces an ε -approximation (2.3) of $\text{vol}(K)$. Then the important result of Dyer, Frieze, and Kanan [47] states that there exists a FPRAS for the volume of a convex body given by a well-guaranteed membership oracle.

Exact methods

Basically, methods for *exact* computation of the volume (triangulations or simplicial decompositions) can be classified according to whether one has a *half-space* description

$$\Omega = \{x \in \mathbb{R}^n | Ax \leq y\}$$

or a *vertex* description

$$\Omega = \left\{ x \in \mathbb{R}^n | x = \sum_{j=1}^p \lambda_j x(j), \lambda_j \geq 0, \sum_{j=1}^p \lambda_j \leq 1 \right\}$$

of Ω , or when both descriptions are available. For instance, Delaunay's triangulation (see, e.g., [29]) and von Hohenbalken's simplicial decomposition [131] both require the list of vertices, whereas Lasserre's algorithm [94] requires a half-space description. On the other hand, Lawrence's [102] and Brion and Vergne's [27] formulas, as well as Cohen and Hickey's triangulation method [34], require both half-space and vertex descriptions of the polytope. On the other hand, Barvinok's algorithm [13] computes the volume by computing the integral of $e^{c'x}$ over Ω for a small c , i.e., evaluates the Laplace transform of the function $x \mapsto I_\Omega(x)$ at the particular $\lambda = c$. We call these approaches *primal* because they all work directly in the primal space \mathbb{R}^n of the variables describing the polytope regardless of whether Ω has a vertex or half-space description.

In this chapter, we take a *dual*¹ approach, that is, we consider problem **I** as that of evaluating the value function $\hat{f}(y, 0)$ at some particular $y \in \mathbb{R}^m$, and to do so, we compute the inverse Laplace transform of its Laplace transform $\hat{F}(\lambda, 0)$ at the point $y \in \mathbb{R}^m$; we call this the *dual integration problem I**. In the present context, as \hat{F} is available in closed form, computing the inverse Laplace transform at the point y is simply the evaluation of a complex integral. This method is *dual* in nature because contrast to primal methods which work in the primal space \mathbb{R}^n of the x variables, we instead work in the space \mathbb{R}^m of *dual* variables λ associated with the nontrivial constraints $Ax = y$. In summary:

- Triangulations or signed decomposition methods, as well as Barvinok's algorithm, are primal methods that work in the space \mathbb{R}^n of primal variables, regardless of whether Ω has a vertex or half-space description. The right-hand side y is fixed.
- Our dual type algorithm works in \mathbb{C}^m as it uses residue techniques to *invert* the Laplace transform $\hat{F}(\lambda, 0) : \mathbb{C}^m \rightarrow \mathbb{C}$ of the function $f(\cdot, 0) : \mathbb{R}^m \rightarrow \mathbb{R}$, at the particular point $y \in \mathbb{R}^m$. So $c = 0$ is fixed, y varies as the argument of $\hat{f}(\cdot, 0)$, and λ varies in the inverse Laplace integral.

2.2 Primal methods

Primal methods for exact volume computation can be divided into two classes: triangulations and signed decompositions. A basic result used in all these methods is the exact formula for the volume of a simplex in \mathbb{R}^n . Let $\Delta(x_0, \dots, x_n)$ be the simplex of \mathbb{R}^n with vertices $x_0, \dots, x_n \in \mathbb{R}^n$. Then

$$\text{vol}(\Delta(x_0, \dots, x_n)) = \frac{|\det(x_1 - x_0, x_2 - x_0, \dots, x_n - x_0)|}{n!}. \quad (2.4)$$

¹ Duality here has nothing to do with duality between the *vertex* and *half-space* descriptions of a convex polytope $\Omega \subset \mathbb{R}^n$.

Triangulations

A *triangulation* of an n -polytope $\Omega \subset \mathbb{R}^n$ is a set $\{\Delta_i\}_{i=1}^s$ of n -simplices $\Delta_i \subset \mathbb{R}^n$ such that no two distinct simplices have an interior point in common. And so,

$$\text{vol}(\Omega) = \sum_{i=1}^s \text{vol}(\Delta_i),$$

with $\text{vol}(\Delta_i)$ given by (2.4). Hence, most of the work is concentrated in finding a triangulation of Ω . For instance, let $\Omega \subset \mathbb{R}^2$ be the rectangle $ABDC$ in Figure 2.1. With E inside Ω , one has the triangulation $(AEC), (AEB), (BED), (DCEC)$ of Ω . Among the triangulation methods are Delaunay's triangulation, boundary triangulation, and Cohen and Hickey's triangulation [34]. While the first two only require a vertex representation of Ω , the latter requires both vertex and half-space descriptions.




FIGURE 2.1 Triangulation of a rectangle

Delaunay's triangulation is obtained by lifting the polytope Ω onto the surface of an $(n+1)$ -dimensional convex body, by $x \mapsto (x, f(x))$ for some strictly convex function $f: \mathbb{R}^n \rightarrow \mathbb{R}$. The lifted vertices $(v, f(v))$ must be in the general position (i.e., their convex hull simplicial). Then, interpreting the facets in terms of the original vertices yields the desired triangulation.

Boundary triangulation for simplicial polytopes (i.e., each facet is an $(n-1)$ -dimensional simplex) links each facet with an interior point as in Figure 2.1 so as to yield the desired triangulations. For general polytopes, a small perturbation yields a simplicial polytope. Cohen and Hickey's triangulation is an improvement of boundary triangulation as it considers a vertex instead of an interior point.

Signed decompositions

As another primal method, one may decompose Ω into *signed* n -simplices $\{\Delta_i\}$ whose signed union is exactly Ω . That is, one writes $\Omega = \bigcup_i \varepsilon_i \Delta_i$ with $\varepsilon_i \in \{-1, +1\}$, and where the signed decomposition $\{\varepsilon_i \Delta_i\}$ must satisfy the following requirement for every point x not on the boundary

of any simplex: If $x \in \Omega$ then x appears in exactly one more positive Δ_i , whereas if $x \notin \Omega$ then it must appear equally often in positive and negative Δ_i . And so, we now have

$$\text{vol}(\Omega) = \sum_{i=1}^s \varepsilon_i \text{vol}(\Delta_i).$$

In Figure 2.2 the same rectangle Ω has the signed decomposition

$$[(AB, AC, e), (e, CD, BD), -(AB, e, BD), -(AC, CD, e)],$$




FIGURE 2.2 Signed decomposition of a rectangle

where a triplet (a, b, c) stands for the unique simplex determined by the intersection of lines a, b , and c . Filliman's duality [56] shows that triangulations and signed decompositions are in duality via the facet–vertex duality (i.e., to a vertex corresponds a facet and viceversa). Namely, let $\Omega \subset \mathbb{R}^n$ be a polytope with the origin in its interior and let T be a triangulation of Ω such that the origin is not in the union of hyperplanes defining the triangulation T . Then T induces a signed decomposition of the polar

$$\Omega^\circ = \{x \in \mathbb{R}^n \mid \langle x, y \rangle \leq 1, \forall y \in \Omega\}$$

of Ω , in such a way that each simplex $\Delta(a_0, \dots, a_n)$ is in T if and only if the unique bounded simplex determined by the corresponding hyperplanes in the polar Ω° is in the signed decomposition. The sign of a simplex in the signed decomposition is determined by its separation parity in the triangulation.

Among signed decomposition methods, let us cite Lawrence's decomposition [102], which results in a simple and elegant formula. With $A \in \mathbb{R}^{m \times n}$, let $\Omega := \{x \in \mathbb{R}^n \mid Ax \leq y\}$ be simple, and for a vertex v of Ω , let $A_v \in \mathbb{R}^{n \times n}$ be the submatrix of A formed with the rows defining the binding constraints at v . Let $c \in \mathbb{R}^n$ and $q \in \mathbb{R}$ be such that $c'x + q$ is not constant along any edge, and let $\pi^v := c'A_v^{-1}$ be well defined. Then

$$\text{vol}(\Omega) = \frac{(c'x + q)^n}{n! |\det A_v| \prod_{i=1}^n \pi_i^v}. \quad (2.5)$$

We will see in Section 2.3 that (2.5) is the same as (2.16) obtained from Brion and Vergne's continuous formula (2.15) (given for Ω described in standard equality form (2.2)).

Alternatively, let $\text{vol}_i(n-1, \Omega_i)$ be the $(n-1)$ -dimensional volume of the facet $\Omega_i := \Omega \cap \{x : (Ax)_i = y_i\}$. Then the formula

$$n \text{vol}(\Omega) = \sum_{i=1}^m \frac{y_i}{\|A_i\|} \text{vol}(n-1, \Omega_i), \quad (2.6)$$

obtained in Lasserre [94] as just Euler's identity for homogeneous functions (applied to $\text{vol}(\Omega)$ as a function of the right-hand side y), has the well-known equivalent geometric formulation

$$n \text{vol}(\Omega) = \sum_{i=1}^m d(O, H_i) \text{vol}(n-1, \Omega_i)$$

(where $d(O, H_i)$ is the algebraic distance from the origin to the hyperplane $A_i x = y_i$), the n -dimensional analogue of the two-dimensional formula (height)(base)/2 for the triangle area. From this geometric interpretation, (2.6) is also a signed decomposition of Ω when O is outside Ω , and a triangulation when O is inside Ω . For instance, if E in Figure 2.1 is the origin and Ω_i is the segment (CD) , then $y_i/\|A_i\|$ is the distance from E to (CD) and $y_i \text{vol}(1, \Omega_i)/2\|A_i\|$ is just the classical formula for the area of the triangle ECD . Formula (2.6) permits us to derive a recursive algorithm as one iterates and one applies again this formula for the polytope Ω_i , for all $i=1, \dots, m$ (see Lasserre [94]).

Concerning the computational complexity and efficiency of triangulation and signed decomposition methods, the interested reader is referred to Büeler et al. [29]. In their computational study, they have observed that triangulation methods behave well for near-simplicial polytopes, whereas signed decomposition methods behave well for near-simple polytopes. This opposite behavior is typical for the hypercube and its polar, the cross polytope. Also, Delaunay's triangulation as well as Lawrence's formula exhibited serious numerical instability in many examples. The latter is due to difficulties in finding a *good* vector c .

Barvinok's algorithm

For each vertex v of $\Omega \subset \mathbb{R}^n$ let $K_v := \text{co}(\Omega, v)$ be the tangent cone (or supporting cone) of Ω at v , that is,

$$\text{co}(\Omega, v) = \{x \in \mathbb{R}^n \mid \varepsilon x + (1 - \varepsilon)v \in \Omega, \forall \text{ sufficiently small } \varepsilon > 0\}. \quad (2.7)$$

Then Brion [25] proved that

$$s(\Omega) = \int_{\Omega} e^{c'x} dx = \sum_{v: \text{vertex of } \Omega} \int_{K_v} e^{c'x} dx =: \sum_{v: \text{vertex of } \Omega} s(K_v).$$

In Barvinok's algorithm [13], each cone K_v is the union $\bigcup_j K_{vj}$ of *simple* cones $\{K_{vj}\}$ so that $s(K_v) = \sum_j s(K_{vj})$. In addition, let K_v be given as the convex hull of its extreme rays $\{u_1, \dots, u_m\} \subset \mathbb{R}^n$. Then

given $c \in \mathbb{C}^m$, one can compute $s(K_v)$ in $O(n^3 m \binom{m}{n})$ arithmetic operations. Finally, taking c very small yields an approximation of the volume of Ω .

As already mentioned, the above methods are primal as they all work in the primal space \mathbb{R}^n of primal variables, regardless of whether Ω has a vertex or half-space representation. We next consider a method that works in the dual space \mathbb{R}^m of variables associated with the nontrivial constraints $Ax = y$ in the description (2.2) of $\Omega(y)$.

2.3 A dual approach

Let $y \in \mathbb{R}^m$ and $A \in \mathbb{R}^{m \times n}$ be real-valued matrices and consider the convex polyhedron $\Omega(y)$ defined in (2.2). Assume that

$$\{x \in \mathbb{R}^n | Ax = 0; c'x \geq 0, x \geq 0\} = \{0\}. \quad (2.8)$$

That is, infinite lines in $\Omega(y)$ with $c'x \geq 0$ are excluded, which makes sense as we want to integrate $e^{c'x}$ over $\Omega(y)$.

Remark 2.1. By a special version of Farkas' lemma due to Carver (see, e.g., Schrijver [121, (33), p. 95]) and adapted to the present context, (2.8) holds if and only if there exists $u \in \mathbb{R}^m$ such that $A'u > c$.

The Laplace transform

Consider the function $\widehat{f}(\cdot, c) : \mathbb{R}^m \rightarrow \mathbb{R}$ defined by

$$y \mapsto \widehat{f}(y, c) := \int_{\Omega(y)} e^{c'x} d\sigma, \quad (2.9)$$

where $d\sigma$ is the Lebesgue measure on the smallest affine variety that contains $\Omega(y)$.

Let $\widehat{F}(\cdot, c) : \mathbb{C}^m \rightarrow \mathbb{C}$ be its m -dimensional two-sided Laplace transform, i.e.,

$$\lambda \mapsto \widehat{F}(\lambda, c) := \int_{\mathbb{R}^m} e^{-\langle \lambda, y \rangle} \widehat{f}(y, c) dy \quad (2.10)$$

(see, e.g., [28], and Appendix).

Theorem 2.2. *Let $\Omega(y)$ be the convex polytope in (2.2) and let \widehat{f} and \widehat{F} be as in (2.9) and (2.10), respectively. Then*

$$\widehat{F}(\lambda, c) = \frac{1}{\prod_{j=1}^n (A'\lambda - c)_j}, \quad \Re(A'\lambda - c) > 0. \quad (2.11)$$

Moreover,

$$\widehat{f}(y, c) = \frac{1}{(2\pi i)^m} \int_{\gamma_1 - i\infty}^{\gamma_1 + i\infty} \cdots \int_{\gamma_m - i\infty}^{\gamma_m + i\infty} e^{\langle \lambda, y \rangle} \widehat{F}(\lambda) d\lambda \quad (2.12)$$

where $\gamma \in \mathbb{R}^m$ satisfies $A'\gamma > c$.

Proof. Apply the definition (2.10) of \widehat{F} , to obtain

$$\begin{aligned} \widehat{F}(\lambda) &= \int_{\mathbb{R}^m} e^{-\langle \lambda, y \rangle} \left[\int_{x \geq 0, Ax=y} e^{c'x} d\sigma \right] dy \\ &= \int_{\mathbb{R}_+^n} e^{\langle c - A'\lambda, x \rangle} dx \\ &= \prod_{j=1}^n \left[\int_0^\infty e^{(c - A'\lambda)_j x_j} dx_j \right] \\ &= \frac{1}{\prod_{j=1}^n (A'\lambda - c)_j}, \text{ with } \Re(A'\lambda - c) > 0. \end{aligned}$$

And (2.12) is obtained by a direct application of the inverse Laplace transform; see [28] and Appendix. It remains to show that the domain $\{\lambda \in \mathbb{C}^m : \Re(A'\lambda - c) > 0\}$ is nonempty. However, this follows from Remark 2.1. \square

The dual problem \mathbf{I}^*

We define the dual problem \mathbf{I}^* to be the *inversion* problem:

$$\begin{aligned} \mathbf{I}^* : \widehat{f}(y, c) &:= \frac{1}{(2i\pi)^m} \int_{\gamma_m - i\infty}^{\gamma_m + i\infty} \cdots \int_{\gamma_1 - i\infty}^{\gamma_1 + i\infty} e^{\langle y, \lambda \rangle} \widehat{F}(\lambda, c) d\lambda \\ &= \frac{1}{(2i\pi)^m} \int_{\gamma_m - i\infty}^{\gamma_m + i\infty} \cdots \int_{\gamma_1 - i\infty}^{\gamma_1 + i\infty} \frac{e^{\langle y, \lambda \rangle}}{\prod_{j=1}^n (A'\lambda - c)_j} d\lambda, \end{aligned} \quad (2.13)$$

where $\gamma \in \mathbb{R}^m$ is fixed and satisfies $A'\gamma - c > 0$.

We may indeed call \mathbf{I}^* a dual problem of \mathbf{I} as it is the inversion of the Laplace transform of \widehat{f} , exactly like \mathbf{P}^* is the inversion of the Legendre–Fenchel transform of f . It is defined on the space \mathbb{C}^m of variables $\{\lambda_j\}$ associated with the nontrivial constraints $Ax = y$; notice that we also retrieve the standard “ingredients” of the dual optimization problem \mathbf{P}^* , namely, $y'\lambda$ and $A'\lambda - c$. Incidentally, observe that the domain of definition (3.7) of $\widehat{F}(\cdot, c)$ is precisely the interior of the feasible set of

the dual problem \mathbf{P}^* in (3.5). We will comment more on this when we compare problems \mathbf{P} and \mathbf{I} in Section 3.2.

As mentioned earlier, computing $\widehat{f}(y, c)$ at the point $y \in \mathbb{R}^m$ via (2.12), i.e., by solving the dual problem \mathbf{I}^* (2.13), reduces to computing the Laplace inverse of $\widehat{F}(\lambda, c)$ at the point y . For instance, this complex integral can be evaluated by repeated application of Cauchy residue theorem. That is, one computes the integral (2.12) by successive one-dimensional complex integrals with respect to (w.r.t.) one variable λ_k at a time (e.g., starting with $\lambda_1, \lambda_2, \dots$) and by repeated application of Cauchy residue theorem. This is possible because the integrand is a rational fraction, and after application of Cauchy residue theorem at step k w.r.t. λ_k , the output is still a rational fraction of the remaining variables $\lambda_{k+1}, \dots, \lambda_m$; see Section 2.4. It is not difficult to see that the whole procedure is a summation of partial results, each of them corresponding to a (multipole) vector $\widehat{\lambda} \in \mathbb{R}^m$ that annihilates m terms of n products in the denominator of the integrand. This is formalized in the nice formula of Brion and Vergne [27, Proposition 3.3, p. 820], which we describe below.

Brion and Vergne's continuous formula

A remarkable result of Brion and Vergne provides us with an *explicit* expression of $\widehat{f}(y, c)$. However, direct evaluation of this formula may pose some numerical instability, and in Section 4.4 we also provide a residue algorithm to evaluate $\widehat{f}(y, c)$ at a point $y \in \mathbb{R}^m$.

Before stating Brion and Vergne's formula, we first introduce some notation. Write the matrix $A \in \mathbb{R}^{m \times n}$ as $A = [A_1 | \dots | A_n]$ where $A_j \in \mathbb{R}^m$ denotes the j th column of A for all $j = 1, \dots, n$. With $\Delta := (A_1, \dots, A_n)$ let $C(\Delta) \subset \mathbb{R}^m$ be the closed convex cone generated by Δ .

A subset σ of $\{1, \dots, n\}$ is called a *basis* of Δ if the sequence $\{A_j\}_{j \in \sigma}$ is a basis of \mathbb{R}^m , and the set of bases of Δ is denoted by $\mathcal{B}(\Delta)$. For $\sigma \in \mathcal{B}(\Delta)$ let $C(\sigma)$ be the cone generated by $\{A_j\}_{j \in \sigma}$. With any $y \in C(\Delta)$ associate the intersection of all cones $C(\sigma)$ that contain y . This defines a subdivision of $C(\Delta)$ into polyhedral cones. The interiors of the maximal cones in this subdivision are called *chambers*; see, e.g., Brion and Vergne [27]. For every $y \in \gamma$, the convex polyhedron $\Omega(y)$ in (2.2) is *simple*. Next, for a chamber γ (whose closure is denoted by $\overline{\gamma}$), let $\mathcal{B}(\Delta, \gamma)$ be the set of bases σ such that γ is contained in $C(\sigma)$, and let $\mu(\sigma)$ denote the volume of the convex polytope $\{\sum_{j \in \sigma} t_j A_j \mid 0 \leq t_j \leq 1\}$ (that is, the determinant of the square matrix $[A_j]_{j \in \sigma}$). Observe that for $y \in \overline{\gamma}$ and $\sigma \in \mathcal{B}(\Delta, \gamma)$ we have $y = \sum_{j \in \sigma} x_j(\sigma) A_j$ for some $x_j(\sigma) \geq 0$. Therefore, the vector $x(\sigma) \in \mathbb{R}_+^n$ with $x_j(\sigma) = 0$ whenever $j \notin \sigma$ is a *vertex* of the polytope $\Omega(y)$. In the LP terminology, the bases $\sigma \in \mathcal{B}(\Delta, \gamma)$ correspond to *feasible bases* of the linear program \mathbf{P} . Denote by V the subspace $\{x \in \mathbb{R}^n \mid Ax = 0\}$. Finally, given $\sigma \in \mathcal{B}(\Delta)$, let $\pi_\sigma \in \mathbb{R}^m$ be the row vector that solves $\pi_\sigma A_j = c_j$ for all $j \in \sigma$. A vector $c \in \mathbb{R}^n$ is said to be *regular* if $c_j - \pi_\sigma A_j \neq 0$ for all $\sigma \in \mathcal{B}(\Delta)$ and all $j \notin \sigma$.

Let $c \in \mathbb{R}^n$ be regular with $-c$ in the interior of the dual cone $(\mathbb{R}_+^n \cap V)^*$ (which is the case if $A'u > c$ for some $u \in \mathbb{R}^m$). Then, Brion and Vergne's formula [27, Proposition 3.3, p. 820] states that

$$\widehat{f}(y, c) = \sum_{\sigma \in \mathcal{B}(\Delta, \gamma)} \frac{e^{\langle c, x(\sigma) \rangle}}{\mu(\sigma) \prod_{k \notin \sigma} (-c_k + \pi_\sigma A_k)} \quad \forall y \in \overline{\gamma}. \quad (2.14)$$

Notice that in the linear programming terminology, $c_k - \pi_\sigma A_k$ is nothing less than the so-called *reduced cost* of the variable x_k , w.r.t. the basis $\{A_j\}_{j \in \sigma}$. When $y \in \gamma$, the polytope $\Omega(y)$ is *simple* and we can rewrite (2.14) as

$$\widehat{f}(y, c) = \sum_{x(\sigma): \text{ vertex of } \Omega(y)} \frac{e^{\langle c, x(\sigma) \rangle}}{\mu(\sigma) \prod_{j \notin \sigma} (-c_j + \pi_\sigma A_j)}. \quad (2.15)$$

On the other hand, when $y \in \bar{\gamma} \setminus \gamma$ for several chambers γ , then in the summation (2.14), several bases σ correspond to the same vertex $x(\sigma)$ of $\Omega(y)$.

Remarkably, evaluating $\widehat{f}(\cdot, c)$ at a point $y \in \mathbb{R}^m$ reduces to sum up over all vertices $x(\sigma)$ of $\Omega(y)$, the quantity $e^{c'x(\sigma)} B_\sigma^{-1}$, where B_σ is just the product of all reduced costs multiplied by the determinant of the matrix basis A_σ . However, this formula is correct only when c is regular. So, for instance, if one wishes to compute the volume $\widehat{f}(y, 0)$ of $\Omega(y)$, i.e., with $c = 0$, one cannot apply directly formula (2.15). However, one also has

$$\text{vol}(\Omega(y)) = \frac{\langle -c, x(\sigma) \rangle^n}{n! \mu(\sigma) \prod_{j \notin \sigma} (c_j - \pi_\sigma A_j)}. \quad (2.16)$$

In fact, (2.16) is exactly Lawrence's formula (2.5) when the polytope $\Omega(y)$ is in standard equality form (2.2). As mentioned already, this formula is not stable numerically, because despite a regular vector $c \in \mathbb{R}^n$ being generic, in many examples the product in the denominator can be very small and one ends up adding and subtracting huge numbers; see, e.g., the empirical study in Büeler et al. [29].

2.4 A residue algorithm for problem I*

We have just seen that evaluating $\widehat{f}(\cdot, c)$ at a point $y \in \mathbb{R}^m$ reduces to evaluating formula (2.15). However, this formula is correct only when c is regular. Moreover, even if c is regular it may pose some numerical difficulties. The same is true for the volume formula (2.16). In this section we provide a residue algorithm that handles the case where c is not regular; when $c = 0$ the algorithm returns the volume of the convex polytope $\Omega(y)$.

As computing the volume of a convex polyhedron is a basic problem of interest in computational geometry, we will specify the algorithm for the case $c = 0$ and for the inequality case description of $\Omega(y)$, i.e., when

$$\Omega(y) := \{x \in \mathbb{R}^n \mid Ax \leq y, x \geq 0\},$$

in which case an easy calculation yields

$$\widehat{F}(\lambda, c) = \frac{1}{\prod_{k=1}^m \lambda_k \prod_{j=1}^n (A' \lambda - c)_j}, \quad \Re(\lambda), \Re(A' \lambda - c) > 0. \quad (2.17)$$

But it will become apparent that the case $c \neq 0$ can be handled exactly in the same manner. Therefore, we wish to evaluate the complex integral

$$\text{volume}(\Omega(y)) = \frac{1}{(2\pi i)^m} \int_{\gamma_1 - i\infty}^{\gamma_1 + i\infty} \cdots \int_{\gamma_m - i\infty}^{\gamma_m + i\infty} e^{y' \lambda} \widehat{F}(\lambda, 0) d\lambda,$$

with $\widehat{F}(y, c)$ as in (2.17) and where the real vector $0 < \gamma \in \mathbb{R}^m$ satisfies $A' \gamma > 0$. To make calculations a little easier, with $y \in \mathbb{R}^m$ fixed, we restate the problem as that of computing the value $h(1)$ of the function $h : \mathbb{R} \rightarrow \mathbb{R}$ given by

$$z \mapsto h(z) := \widehat{f}(zy, 0) = \frac{1}{(2\pi i)^m} \int_{\gamma_1 - i\infty}^{\gamma_1 + i\infty} \cdots \int_{\gamma_m - i\infty}^{\gamma_m + i\infty} e^{z \langle \lambda, y \rangle} \widehat{F}(\lambda, 0) d\lambda. \quad (2.18)$$

Computing the complex integral (2.18) can be done in two ways that are explored below. We do it directly in Section 2.4 by integrating first w.r.t. λ_1 , then w.r.t. $\lambda_2, \lambda_3, \dots$; or indirectly in Section 2.4, by first computing the one-dimensional Laplace transform H of h and then computing the Laplace inverse of H .

The direct method

To better understand the methodology behind the direct method and for illustration purposes, consider the case of a convex polytope $\Omega(y)$ with only two ($m = 2$) nontrivial constraints.

The $m = 2$ example. Let $A \in \mathbb{R}^{2 \times n}$ be such that $x = 0$ is the only solution of $\{x \geq 0, Ax \leq 0\}$ so that $\Omega(y)$ is a convex polytope (see Remark 2.1). Moreover, suppose that $y_1 = y_2 = z$ and $A' := [a \mid b]$ with $a, b \in \mathbb{R}^n$. For ease of exposition, assume that

- $a_j b_j \neq 0$ and $a_j \neq b_j$ for all $j = 1, 2, \dots, n$.
- $a_j / b_j \neq a_k / b_k$ for all $j, k = 1, 2, \dots, n$.

Observe that these assumptions are satisfied with probability equal to one if we add to every coefficient a_i, b_i a perturbation ε , randomly generated under a uniform distribution on $[0, \bar{\varepsilon}]$, with $\bar{\varepsilon}$ very small. Then

$$\widehat{F}(\lambda, 0) = \frac{1}{\lambda_1 \lambda_2} \times \frac{1}{\prod_{j=1}^n (a_j \lambda_1 + b_j \lambda_2)}, \quad \begin{cases} \Re(\lambda) > 0 \\ \Re(a \lambda_1 + b \lambda_2) > 0. \end{cases}$$

Next, fix γ_1 and $\gamma_2 > 0$ such that $a_j \gamma_1 + b_j \gamma_2 > 0$ for every $j = 1, 2, \dots, n$ and compute the integral (2.18) as follows. We first evaluate the integral

$$I_1(\lambda_2) := \frac{1}{2\pi i} \int_{\gamma_1 - i\infty}^{\gamma_1 + i\infty} \frac{e^{z \lambda_1}}{\lambda_1 \prod_{j=1}^n (a_j \lambda_1 + b_j \lambda_2)} d\lambda_1, \quad (2.19)$$

using classical Cauchy residue technique. That is, we (a) close the path of integration adding a semicircle Γ of radius R large enough, (b) evaluate the integral using Cauchy residue theorem (see Appendix), and (c) show that the integral along Γ converges to zero when $R \rightarrow \infty$.

Now, since we are integrating w.r.t. λ_1 and we want to evaluate $h(z)$ at $z = 1$, we must add the semicircle Γ on the left side of the integration path $\{\Re(\lambda_1) = \gamma_1\}$ because $e^{z\lambda_1}$ converges to zero when $\Re(\lambda_1) \rightarrow -\infty$. Therefore, we must consider *only* the poles of $\lambda_1 \mapsto \hat{F}(\lambda_1, \lambda_2)$ whose real part is strictly less than γ_1 (with λ_2 being fixed). Recall that $\Re(-\lambda_2 b_j/a_j) = -\gamma_2 b_j/a_j < \gamma_1$ whenever $a_j > 0$, for each $j = 1, 2, \dots, n$, and $\hat{F}(\lambda_1, \lambda_2)$ has only poles of first order (with λ_2 being fixed). Then, the evaluation of (2.19) follows easily, and

$$I_1(\lambda_2) = \frac{1}{\lambda_2^n \prod_{j=1}^n b_j} + \sum_{a_j > 0} \frac{-e^{-(b_j/a_j)z\lambda_2}}{b_j \lambda_2^n \prod_{k \neq j} (-a_k b_j/a_j + b_k)}.$$

Therefore,

$$\begin{aligned} h(z) &= \frac{1}{2\pi i} \int_{\gamma_2 - i\infty}^{\gamma_2 + i\infty} \frac{e^{z\lambda_2}}{\lambda_2} I_1(\lambda_2) d\lambda_2 \\ &= \frac{1}{2\pi i} \int_{\gamma_2 - i\infty}^{\gamma_2 + i\infty} \frac{e^{z\lambda_2}}{\lambda_2^{n+1} \prod_{j=1}^n b_j} d\lambda_2 \\ &\quad - \sum_{a_j > 0} \frac{1}{2\pi i} \int_{\gamma_2 - i\infty}^{\gamma_2 + i\infty} \frac{a_j^n e^{(1-b_j/a_j)z\lambda_2}}{\lambda_2^{n+1} a_j b_j \prod_{k \neq j} (b_k a_j - a_k b_j)} d\lambda_2. \end{aligned}$$

These integrals must be evaluated according to whether $(1 - b_j/a_j)y$ is positive or negative, or, equivalently (as $a_j > 0$), whether $a_j > b_j$ or $a_j < b_j$. Thus, recalling that $z > 0$, each integral is equal to

- its residue at the pole $\lambda_2 = 0 < \gamma_2$ when $a_j > b_j$, and
- zero if $a_j < b_j$ because there is no pole on the right side of $\{\Re(\lambda_2) = \gamma_2\}$.

That is,

$$h(z) = \frac{z^n}{n!} \left[\frac{1}{\prod_{j=1}^n b_j} - \sum_{b_j < a_j} \frac{(a_j - b_j)^n}{a_j b_j \prod_{k \neq j} (b_k a_j - a_k b_j)} \right]. \quad (2.20)$$

Observe that the formula is not symmetric in the parameters a, b . This is because we have chosen to integrate first w.r.t. λ_1 ; and the set $\{j | b_j < a_j\}$ is different from $\{j | a_j < b_j\}$, which would have been considered had we integrated first w.r.t. λ_2 . In the latter case, we would have obtained

$$h(z) = \frac{z^n}{n!} \left[\frac{1}{\prod_{j=1}^n a_j} - \sum_{a_j < b_j} \frac{(b_j - a_j)^n}{a_j b_j \prod_{k \neq j} (a_k b_j - b_k a_j)} \right], \quad (2.21)$$

which is (2.20), interchanging a and b . Moreover, moving terms around, we get for free the following identity:

$$\sum_{j=1}^n \frac{(a_j - b_j)^n}{a_j b_j \prod_{k \neq j} (b_k a_j - a_k b_j)} = \frac{1}{\prod_{j=1}^n b_j} - \frac{1}{\prod_{j=1}^n a_j}. \quad (2.22)$$

The direct method algorithm. The above methodology easily extends to an arbitrary number m of nontrivial constraints. One evaluates the integral of the right hand side of (2.18) by integrating first w.r.t. λ_1 , then w.r.t. λ_2 , and so on. The resulting algorithm can be described with a tree of depth $m+1$ ($m+1$ “levels”).

Direct method algorithm: Let $0 < c \in \mathbb{R}^m$ be such that $A'c > 0$.

- Level 0 is the root of the tree.
 - Level 1 is the integration w.r.t. λ_1 and consists of at most $n+1$ nodes associated with the poles $\lambda_1 := \rho_j^1$, $j = 1, \dots, n+1$, of the rational function $\prod_i \lambda_i^{-1} \prod_j (A'\lambda)_j^{-1}$, seen as a function of λ_1 only. By the assumption on c , there is no pole ρ_j^1 on the line $\Re(\lambda_1) = \gamma_1$. By the Cauchy residue theorem, only the poles at the left side of the integration path $\{\Re(\lambda_1) = \gamma_1\}$, say ρ_j^1 , $j \in I_1$, are selected.
 - Level 2 is the integration w.r.t. λ_2 . After integration w.r.t. λ_1 , each of the poles ρ_j^1 , $j \in I_1$, generates a rational function of $\lambda_2, \lambda_3, \dots, \lambda_m$ times an exponential, which, seen as a function of λ_2 only, has at most $n+1$ poles $\rho_i^2(j)$, $i = 1, \dots, n+1$, $j \in I_1$. Thus, level 2 has at most $(n+1)^2$ nodes associated with the poles $\rho_i^2(j)$. Assuming no pole $\rho_i^2(j)$ on the line $\Re(\lambda_2) = \gamma_2$, by the Cauchy residue theorem, only the poles $\rho_i^2(j)$, $(j, i) \in I_2$, located on the correct side of the integration path $\{\Re(\lambda_2) = \gamma_2\}$ are selected. By “correct” side we mean the right side if the coefficient in the exponential is negative, and the left side otherwise.
 - Level k , $k \leq m$, consists of at most $(n+1)^k$ nodes associated with the poles $\rho_s^k(i_1, i_2, \dots, i_{k-1})$, $(i_1, i_2, \dots, i_{k-1}) \in I_{k-1}$, $s = 1, \dots, n+1$, of some rational functions of $\lambda_k, \dots, \lambda_m$, seen as functions of λ_k only. Assuming no pole on the line $\Re(\lambda_k) = \gamma_k$, only the poles $\rho_{i_k}^k(i_1, i_2, \dots, i_{k-1})$, $(i_1, i_2, \dots, i_k) \in I_k$, located on the correct side of the integration path $\{\Re(\lambda_k) = \gamma_k\}$ are selected. And so on.
- The last level m consists of at most $(n+1)^m$ nodes and the integration w.r.t. λ_m is trivial as it amounts to evaluating integrals of the form

$$(2\pi i)^{-1} \int_{\gamma_m - i\infty}^{\gamma_m + i\infty} A \lambda_m^{-(n+1)} e^{\alpha z \lambda_m} d\lambda_m$$

for some coefficients A , α , which yield $A(\alpha z)^n/n!$ for those $\alpha > 0$. Summing up over all nodes of the last level (analogously, all leaves of the tree) provides the desired value.

Computational complexity. Only simple elementary arithmetic operations are needed to compute the nodes at each level, as in Gauss elimination for solving linear systems. Moreover, finding a vector $\gamma \in \mathbb{R}^m$ that satisfies $\gamma > 0$, $A'\gamma > 0$, can be done in time polynomial in the input size of the problem. Therefore, the computational complexity is easily seen to be essentially described by n^m .

However, some care must be taken with the choice of the integration paths as we assume that at each level k there is no pole on the integration path $\{\Re(\lambda_k) = \gamma_k\}$. This issue is discussed later. The algorithm is illustrated on the following simple example with $n = 2, m = 3$.

Example 2.1. Let $e_3 = (1, 1, 1)$ be the unit vector of \mathbb{R}^3 and let $\Omega(ze_3) \subset \mathbb{R}^2$ be the polytope

$$\Omega(ze_3) := \{x \in \mathbb{R}_+^2 \mid x_1 + x_2 \leq z; -2x_1 + 2x_2 \leq z; 2x_1 - x_2 \leq z\},$$

whose area is $17z^2/48$.

Choose $\gamma_1 = 3$, $\gamma_2 = 2$, and $\gamma_3 = 1$ so that $\gamma_1 > 2\gamma_2 - 2\gamma_3$ and $\gamma_1 > \gamma_3 - 2\gamma_2$.

$$h(z) = \frac{1}{(2\pi i)^3} \int_{\gamma_1 - i\infty}^{\gamma_1 + i\infty} \dots \int_{\gamma_3 - i\infty}^{\gamma_3 + i\infty} e^{(\lambda_1 + \lambda_2 + \lambda_3)z} \widehat{F}(\lambda) d\lambda,$$

with

$$\widehat{F}(\lambda, 0) = \frac{1}{\lambda_1 \lambda_2 \lambda_3 (\lambda_1 - 2\lambda_2 + 2\lambda_3)(\lambda_1 + 2\lambda_2 - \lambda_3)}.$$

Integrate first w.r.t. λ_1 ; that is, evaluate the residues at the poles $\lambda_1 = 0$, $\lambda_1 = 2\lambda_2 - 2\lambda_3$, and $\lambda_1 = \lambda_3 - 2\lambda_2$ because $0 < z$, $0 < \gamma_1$, $\Re(2\lambda_2 - 2\lambda_3) < \gamma_1$, and $\Re(\lambda_3 - 2\lambda_2) < \gamma_1$. We obtain

$$h(z) = \frac{1}{(2\pi i)^2} \int_{\gamma_2 - i\infty}^{\gamma_2 + i\infty} \int_{\gamma_3 - i\infty}^{\gamma_3 + i\infty} (I_2 + I_3 + I_4) d\lambda_2 d\lambda_3,$$

where

$$\begin{aligned} I_2(\lambda_2, \lambda_3) &= \frac{-e^{(\lambda_2 + \lambda_3)z}}{2\lambda_2 \lambda_3 (\lambda_3 - \lambda_2)(\lambda_3 - 2\lambda_2)}, \\ I_3(\lambda_2, \lambda_3) &= \frac{e^{(3\lambda_2 - \lambda_3)z}}{6\lambda_2 \lambda_3 (\lambda_3 - \lambda_2)(\lambda_3 - 4\lambda_2/3)}, \\ I_4(\lambda_2, \lambda_3) &= \frac{e^{(2\lambda_3 - \lambda_2)z}}{3\lambda_2 \lambda_3 (\lambda_3 - 2\lambda_2)(\lambda_3 - 4\lambda_2/3)}. \end{aligned}$$

Next, integrate $I_2(\lambda_2, \lambda_3)$ w.r.t. λ_3 . We must consider the poles on the left side of $\{\Re(\lambda_3) = 1\}$, that is, the pole $\lambda_3 = 0$ because $\Re(\lambda_2) = 2$. Thus, we get $-e^{z\lambda_2}/4\lambda_2^3$, and the next integration w.r.t. λ_2 yields $-z^2/8$.

When integrating $I_3(\lambda_2, \lambda_3)$ w.r.t. λ_3 , we have to consider the poles $\lambda_3 = \lambda_2$ and $\lambda_3 = 4\lambda_2/3$ on the right side of $\{\Re(\lambda_3) = 1\}$; and we get

$$\frac{-1}{\lambda_2^3} \left[-\frac{e^{2z\lambda_2}}{2} + \frac{3e^{z\lambda_2 5/3}}{8} \right].$$

Recall that the path of integration has a negative orientation, so we have to consider the negative value of residues. The next integration w.r.t. λ_2 yields $z^2(1 - 25/48)$.

Finally, when integrating $I_4(\lambda_2, \lambda_3)$ w.r.t. λ_3 , we must consider only the pole $\lambda_3 = 0$, and we get $e^{-z\lambda_2}/8\lambda_2^3$; the next integration w.r.t. λ_2 yields zero. Hence, adding up the above three partial results yields

$$h(z) = z^2 \left[\frac{-1}{8} + 1 - \frac{25}{48} \right] = \frac{17z^2}{48},$$

which is the desired result.

Multiple poles. In the description of the algorithm we have considered the case in which all the poles are *simple* at each step. This is indeed the *generic* case. However, multiple poles may occur at some of the first $m - 1$ one-dimensional integrations. For instance, at the k th step (after integration

w.r.t. $\lambda_1, \lambda_2, \dots, \lambda_{k-1}$), the integrand could have a denominator that includes one (or several) product(s) of the form

$$(a_k \lambda_k + \dots + a_m \lambda_m) \times (b_k \lambda_k + \dots + b_m \lambda_m)$$

with

$$\frac{a_k}{b_k} = \frac{a_{k+1}}{b_{k+1}} = \dots = \frac{a_m}{b_m}; \quad (2.23)$$

i.e., there is a pole of order greater than or equal to 2 w.r.t. λ_k .

This is a *pathological* case, which happens with probability zero on a sample of problems with randomly generated data. Moreover, one may handle this situation in one of two ways:

- *Directly*, by detecting multiple poles and applying an adequate Cauchy residue technique. This procedure requires some extra work (*detection* of (2.23) + *derivation* in Cauchy residue theorem), which is a polynomial in the dimension n . Observe that detection of (2.23) requires some ε -tolerance as it deals with real numbers.
- *Indirectly*, by preprocessing (perturbing) the nonzero elements of the initial matrix A . Let us say $A(i, j) \rightarrow A(i, j) + \varepsilon(i, j)$ with independent identically distributed (i.i.d.) random variables $\{\varepsilon(i, j)\}$ taking values in some interval $[-\bar{\varepsilon}, \bar{\varepsilon}]$ for $\bar{\varepsilon}$ very small.

While the former technique requires extra work, the latter is straightforward, but only provides (with probability equal to one) an approximation of the exact volume.

Paths of integration. In choosing the integration paths $\{\Re(\lambda_k) = \gamma_k\}, k = 1, \dots, m$, we must determine a vector $0 < \gamma \in \mathbb{R}^m$ such that $A'\gamma > 0$. However, this may not be enough when we want to evaluate the integral (2.18) by repeated applications of the Cauchy residue theorem. Indeed, in the tree description of the algorithm, we have seen that at each level $k > 1$ of the tree (integration w.r.t. λ_k), one *assumes* that there is *no* pole on the integration path $\{\Re(\lambda_k) = \gamma_k\}$.

For instance, had we set $\gamma_1 = \gamma_2 = \gamma_3 = 1$ (instead of $\gamma_1 = 3, \gamma_2 = 2$, and $\gamma_3 = 1$) in the above example, we could not use Cauchy residue theorem to integrate I_2 or I_3 because we would have the pole $\lambda_2 = \lambda_3$ exactly on the path of integration (recall that $\Re(\lambda_2) = \Re(\lambda_3) = 1$); fortunately, this case is *pathological* as it happens with probability zero in a set of problems with randomly generated data $A \in \mathbb{R}^{m \times n}$, and therefore this issue could be ignored in practice. However, for the sake of mathematical rigor, in addition to the constraints $\gamma > 0$ and $A'\gamma > 0$, the vector $\gamma \in \mathbb{R}^m$ must satisfy additional constraints to avoid the above-mentioned pathological problem. We next describe one way to proceed to ensure that c satisfies these additional constraints.

We have described the algorithm as a tree of depth m (level i being the integration w.r.t. λ_i), where each node has at most $n + 1$ descendants (one descendant for each pole on the correct side of the integration path $\{\Re(\lambda_i) = \gamma_i\}$). The volume is then the summation of all partial results obtained at each leaf of the tree (that is, each node of level m). We next describe how to “perturb” *on-line* the initial vector $c \in \mathbb{R}^m$ if at some level k of the algorithm there is a pole on the corresponding integration path $\{\Re(\lambda_k) = \gamma_k\}$.

Step 1. Integration w.r.t. λ_1 . Choose a real vector $\gamma := (\gamma_1^1, \dots, \gamma_m^1) > 0$ such that $A'\gamma > 0$ and integrate (2.18) along the line $\Re(\lambda_1) = \gamma_1^1$. From the Cauchy residue theorem, this is done by selecting the (at most $n + 1$) poles $\lambda_1 := \rho_j^1, j \in I_1$, located on the left side of the integration path $\{\Re(\lambda_1) = \gamma_1^1\}$. Each pole $\rho_j^1, j = 1, \dots, n + 1$ (with $\rho_j^1 := 0$) is a linear combination

$\beta_{j_2}^{(1)}\lambda_2 + \cdots + \beta_{j_m}^{(1)}\lambda_m$ with real coefficients $\{\beta_{jk}^{(1)}\}$, because A is a real-valued matrix. Observe that by the initial choice of γ ,

$$\delta_1 := \min_{j=1,\dots,n+1} |\gamma_1^1 - \sum_{k=2}^m \beta_{jk}^{(1)} \gamma_k^1| > 0.$$

Step 2. Integration w.r.t. λ_2 . For each of the poles ρ_j^1 , $j \in I_1$, selected at step 1, and after integration w.r.t. λ_1 , we now have to consider a rational function of λ_2 with at most $n+1$ poles $\lambda_2 := \rho_i^2(j) := \sum_{k=3}^m \beta_{ik}^{(2)}(j) \lambda_k$, $i = 1, \dots, n+1$. If

$$\delta_2 := \min_{j \in I_1} \min_{i=1,\dots,n+1} \left| \gamma_2^1 - \sum_{k=3}^m \beta_{ik}^{(2)}(j) \gamma_k^1 \right| > 0,$$

then integrate w.r.t. λ_2 on the line $\Re(\lambda_2) = \gamma_2^1$. Otherwise, if $\delta_2 = 0$ we set $\gamma_2^2 := \gamma_2^1 + \varepsilon_2$ and $\gamma_k^2 := \gamma_k^1$ for all $k \neq 2$, by choosing $\varepsilon_2 > 0$ small enough to ensure that

- (a) $A' \gamma^2 > 0$
- (b) $\delta_2 := \min_{j \in I_1} \min_{i=1,\dots,n+1} \left| \gamma_2^2 - \sum_{k=3}^m \beta_{ik}^{(2)}(j) \gamma_k^2 \right| > 0$
- (c) $\max_{j=1,\dots,n+1} |\beta_{j2}^{(1)} \varepsilon_2| < \delta_1$.

The condition (a) is basic whereas (b) ensures that there is no pole on the integration path $\{\Re(\lambda_2) = \gamma_2^2\}$. Moreover, what has been done in step 1 remains valid because from (c), $\gamma_1^2 - \sum_{k=2}^m \beta_{jk}^{(1)} \gamma_k^2$ has the same sign as $\gamma_1^1 - \sum_{k=2}^m \beta_{jk}^{(1)} \gamma_k^1$, and, therefore, none of the poles ρ_j^1 has crossed the integration path $\{\Re(\lambda_1) = \gamma_1^1 = \gamma_1^2\}$, that is, the set I_1 is unchanged.

Now integrate w.r.t. λ_2 on the line $\Re(\lambda_2) = \gamma_2^2$, which is done via the Cauchy residue theorem by selecting the (at most $(n+1)^2$) poles $\rho_i^2(j)$, $(j, i) \in I_2$, located at the left or the right of the line $\Re(\lambda_2) = \gamma_2^2$, depending on the sign of the coefficient of the argument in the exponential.

Step 3. Integration w.r.t. λ_3 . Likewise, for each of the poles $\rho_i^2(j)$, $(j, i) \in I_2$, selected at step 2, we now have to consider a rational function of λ_3 with at most $n+1$ poles $\rho_s^3(j, i) := \sum_{k=4}^m \beta_{sk}^{(3)}(j, i) \lambda_k$, $s = 1, \dots, n+1$. If

$$\delta_3 := \min_{(j,i) \in I_2} \min_{s=1,\dots,n+1} \left| \gamma_3^2 - \sum_{k=4}^m \beta_{sk}^{(3)}(j, i) \gamma_k^2 \right| > 0,$$

then integrate w.r.t. λ_3 on the line $\Re(\lambda_3) = \gamma_3^2$. Otherwise, if $\delta_3 = 0$, set $\gamma_3^3 := \gamma_3^2 + \varepsilon_3$ and $\gamma_k^3 := \gamma_k^2$ for all $k \neq 3$ by choosing $\varepsilon_3 > 0$ small enough to ensure that

- (a) $A' \gamma^3 > 0$
- (b) $\delta_3 := \min_{(j,i) \in I_2} \min_{s=1,\dots,n+1} \left| \gamma_3^3 - \sum_{k=4}^m \beta_{sk}^{(3)}(j, i) \gamma_k^3 \right| > 0$
- (c) $\max_{j \in I_1} \max_{i=1,\dots,n+1} |\beta_{i3}^{(2)}(j) \varepsilon_3| < \delta_2$
- (d) $\max_{j=1,\dots,n+1} |\beta_{j2}^{(1)} \varepsilon_2 + \beta_{j3}^{(1)} \varepsilon_3| < \delta_1$.

As in previous steps, condition (a) is basic. The condition (b) ensures that there is no pole on the integration path $\{\Re(\lambda_3) = \gamma_3^3\}$. Condition (c) (resp., (d)) ensures that none of the poles $\rho_i^2(j)$ considered at step 2 (resp., none of the poles ρ_j^1 considered at step 1) has crossed the line $\Re(\lambda_2) = \gamma_2^3 = \gamma_2^2$ (resp., the line $\Re(\lambda_1) = \gamma_1^3 = \gamma_1^1$). That is, both sets I_1 and I_2 are unchanged.

Next, integrate w.r.t. λ_3 on the line $\Re(\lambda_3) = \gamma_3^3$, which is done by selecting the (at most $(n+1)^3$) poles $\rho_s(j, i)$, $(j, i, s) \in I_3$, located at the left or right of the line $\{\Re(\lambda_3) = \gamma_3^3\}$, depending on the sign of the argument in the exponential. And so on.

Remark 2.2. (i) It is important to notice that ε_k and γ_k^k play no (numerical) role in the integration itself. They are only used to (1) ensure the absence of a pole on the integration path $\{\Re(\lambda_k) = \gamma_k^k\}$ and (2) locate the poles on the left or the right of the integration path. Their numerical value (which can be very small) has no influence on the computation.

(ii) If one ignores the perturbation technique for the generic case, then a depth-first search of the tree requires a polynomial storage space, whereas an exponential amount of memory is needed with the perturbation technique.

The associated transform algorithm

We now describe an alternative to the direct method, one that permits us (a) to avoid evaluating integrals of exponential functions in (2.18) and (b) to avoid making the *on-line* changes of the integration paths described earlier for handling possible (pathological) poles on the integration path.

Recall that we want to compute (2.18) where $y \neq 0$ and $\gamma > 0$ are real vectors in \mathbb{R}^m with $A'\gamma > 0$. It is easy to deduce from (2.1) that $\hat{f}(y, 0) = 0$ whenever $y \leq 0$, so we may suppose without loss of generality that $y_m > 0$; we make the following simple change of variables as well. Let $p = \langle \lambda, y \rangle$ and $d = \langle \gamma, y \rangle$, so $\lambda_m = (p - \sum_{j=1}^{m-1} y_j \lambda_j) / y_m$ and

$$h(z) = \frac{1}{(2\pi i)^m} \int_{\gamma_1 - i\infty}^{\gamma_1 + i\infty} \cdots \int_{\gamma_{m-1} - i\infty}^{\gamma_{m-1} + i\infty} \left[\int_{d - i\infty}^{d + i\infty} e^{zp} \widehat{G} dp \right] d\lambda_1 \dots d\lambda_{m-1},$$

where

$$\widehat{G}(\lambda_1, \dots, \lambda_{m-1}, p) = \frac{1}{y_m} \widehat{F} \left(\lambda_1, \dots, \lambda_{m-1}, \frac{p}{y_m} - \sum_{j=1}^{m-1} \frac{y_j}{y_m} \lambda_j \right). \quad (2.24)$$

We can rewrite $h(z)$ as follows:

$$h(z) = \frac{1}{2\pi i} \int_{d - i\infty}^{d + i\infty} e^{zp} H(p) dp, \quad (2.25)$$

with

$$H(p) := \frac{1}{(2\pi i)^{m-1}} \int_{\gamma_1 - i\infty}^{\gamma_1 + i\infty} \cdots \int_{\gamma_{m-1} - i\infty}^{\gamma_{m-1} + i\infty} \widehat{G} d\lambda_1 \dots d\lambda_{m-1}. \quad (2.26)$$

Notice that $H(p)$ is the Laplace transform of $h(z) = \hat{f}(zy, 0)$, called the *associated transform* of $\hat{F}(\lambda)$. In addition, since $h(z)$ is proportional to z^n (volume is proportional to the n th power of length),

$$H(p) = C/p^{n+1},$$

where the constant $C = h(1)n!$ gives us the value $h(1) = \widehat{f}(y, 0)$ we are looking for. From here, we need only to compute $H(p)$ in (2.26) in order to know the final result. We can do so via repeated application of the Cauchy residue theorem (as in the direct method algorithm of Section 2.4).

Moreover, after each one of the successive $m - 1$ one-dimensional integrations calculated in (2.26), we *always* get an *analytic* function. This fact can be proved as follows. Consider for some time the change of variables $s_j = 1/\lambda_j$ for all $j = 1, 2, \dots, m - 1$ in (2.26). This change of variables is well-defined because no λ_j is equal to zero, and we obtain

$$H(p) = \frac{1}{(2\pi i)^{m-1}} \int_{|2\gamma_1 s_1 - 1| = 1} \cdots \int_{|2\gamma_{m-1} s_{m-1} - 1| = 1} R ds_1 \dots ds_{m-1},$$

where $R(s_1, \dots, s_{m-1}, p)$ is a rational function (recall that G and \widehat{G} are both rational) continuous on the Cartesian product of integration paths

$$E = \{|2\gamma_1 s_1 - 1| = 1\} \times \cdots \times \{|2\gamma_{m-1} s_{m-1} - 1| = 1\} \times \{\Re(p) = d\}.$$

Recall that $\widehat{F}(\lambda)$ is defined on the domain $\{\Re(\lambda) > 0, \Re(A'\lambda) > 0\}$, and so R is analytic on the domain $D \subset \mathbb{C}^m$ of all points $(s_1, \dots, s_{m-1}, p) \in \mathbb{C}^m$ such that $|2\beta_j s_j - 1| = 1$, $s_j \neq 0$, $j = 1, 2, \dots, m - 1$, and $\Re(p) = \langle \beta, y \rangle$ for some vector $\beta > 0$ in \mathbb{R}^m with $A'\beta > 0$.

Observe that $E \subset D \cup \{0\}$ and each integration path $\{|2\gamma_j s_j - 1| = 1\}$ is a compact circle, $j = 1, \dots, m - 1$. Therefore, after integration w.r.t. s_1, s_2, \dots, s_{m-1} , we get that $H(p)$ is *analytic* on the domain of all points $p \in \mathbb{C}$ such that $\Re(p) = \langle \beta, b \rangle$ for some real vector $\beta = (\gamma_1, \dots, \gamma_{m-1}, \beta_m) > 0$ with $A'\beta > 0$.

Similarly, if we only integrate w.r.t. s_1, s_2, \dots, s_k , $1 \leq k \leq m - 2$, we get a function which is *analytic* on the domain of all points $(s_{k+1}, \dots, s_{m-1}, p) \in \mathbb{C}^{m-k}$ such that $|2\beta_j s_j - 1| = 1$, $s_j \neq 0$, $j = k + 1, \dots, m - 1$, and $\Re(p) = \langle \beta, y \rangle$ for some real vector $\beta = (\gamma_1, \dots, \gamma_k, \beta_{k+1}, \dots, \beta_m) > 0$ with $A'\beta > 0$.

Thus, there is no pole on the next integration path $\{|2\gamma_{k+1} s_{k+1} - 1| = 1\}$ or the final one $\{\Re(p) = d\}$. Referring back to the original variables $\{\lambda_k\}$ and (2.26), we translate this fact as follows. After integrating (2.26) w.r.t. $\lambda_1, \lambda_2, \dots, \lambda_k$, any *apparent* pole placed on the next integration path $\{\Re(\lambda_{k+1}) = \gamma_{k+1}\}$ or the final one $\{\Re(p) = d\}$ is a *removable singularity* (see [36, p. 103]). In other words, poles on the next integration path can be removed by *adding together all terms with poles there*. The general method is illustrated with the same example of two nontrivial constraints ($m = 2$) already considered at the beginning of Section 2.4.

The $m = 2$ example. Recall that $A' := [a \mid b]$ with $a, b \in \mathbb{R}^n$ so that

$$\widehat{F}(\lambda, 0) = \frac{1}{\lambda_1 \lambda_2} \times \frac{1}{\prod_{j=1}^n (a_j \lambda_1 + b_j \lambda_2)}, \quad \Re(\lambda) > 0, \Re(A'\lambda) > 0$$

and fix $y_1 = y_2 = z$. To compare with the direct method, and as in the beginning of Section 2.4, assume that $a_j b_j \neq 0$ for all $j = 1, \dots, n$ and $a_j/b_j \neq a_k/b_k$ for all $j \neq k$.

Fix $\lambda_2 = p - \lambda_1$ and choose real constants $\gamma_1 > 0$ and $\gamma_2 > 0$ such that $a_j \gamma_1 + b_j \gamma_2 > 0$ for every $j = 1, 2, \dots, n$. Notice that $\Re(p) = \gamma_1 + \gamma_2$. We obtain $H(p)$ by integrating $\widehat{F}(\lambda_1, p - \lambda_1, 0)$ on the line $\Re(\lambda_1) = \gamma_1$, which yields

$$H(p) = \frac{1}{2\pi i} \int_{\gamma_1 - i\infty}^{\gamma_1 + i\infty} \frac{1}{\lambda_1(p - \lambda_1)} \times \frac{1}{\prod_{j=1}^n ((a_j - b_j)\lambda_1 + b_j p)} d\lambda_1.$$

Next, we need to determine which poles of $\widehat{F}(\lambda_1, p - \lambda_1)$ are on the left (right) side of the integration path $\{\Re(\lambda_1) = \gamma_1\}$ in order to apply the Cauchy residue theorem. Let $J_+ = \{j | a_j > b_j\}$, $J_0 = \{j | a_j = b_j\}$, and $J_- = \{j | a_j < b_j\}$. Then, the poles on the left side of $\{\Re(\lambda_1) = \gamma_1\}$ are $\lambda_1 = 0$ and $\lambda_1 = -b_j p / (a_j - b_j)$ for all $j \in J_+$ because $-b_j \Re(p) / (a_j - b_j) < \gamma_1$. Additionally, the poles on the right side of $\{\Re(\lambda_1) = \gamma_1\}$ are $\lambda_1 = p$ and $\lambda_1 = -b_j p / (a_j - b_j)$ for all $j \in J_-$. Finally, notice that $\widehat{F}(\lambda_1, p - \lambda_1, 0)$ has only poles of first order.

Hence, computing the residues of poles on the left side of $\{\Re(\lambda_1) = \gamma_1\}$ yields

$$H(p) = \frac{1}{\prod_{j \in J_0} p b_j} \left[\frac{1}{p \prod_{j \notin J_0} p b_j} + \sum_{j \in J_+} \frac{-(a_j - b_j)^{n-|J_0|}}{p^2 a_j b_j \prod_{k \notin J_0, k \neq j} (-p b_j a_k + p a_j b_k)} \right].$$

Hence, after moving terms around, we obtain

$$H(p) = \frac{1}{p^{n+1}} \left[\frac{1}{\prod_{j=1}^n b_j} - \sum_{a_j > b_j} \frac{(a_j - b_j)^n}{a_j b_j \prod_{k \neq j} (a_j b_k - b_j a_k)} \right].$$

Notice that previous equation holds even for the case $J_0 \neq \emptyset$. Finally, after integration w.r.t. p , we get

$$h(z) = \frac{z}{n!} \left[\frac{1}{\prod_{j=1}^n b_j} - \sum_{a_j > b_j} \frac{(a_j - b_j)^n}{a_j b_j \prod_{k \neq j} (a_j b_k - b_j a_k)} \right],$$

which coincides with (2.20) in the particular case $J_0 = \emptyset$.

Now, computing the negative value of residues of poles on the right side of $\{\Re(\lambda_1) = \gamma_1\}$ (we need to take the negative value because the path of integration has a negative orientation), yields

$$H(p) = \frac{1}{p^{n+1}} \left[\frac{1}{\prod_{j=1}^n a_j} - \sum_{b_j > a_j} \frac{(b_j - a_j)^n}{a_j b_j \prod_{k \neq j} (b_j a_k - a_j b_k)} \right],$$

and after integration w.r.t. p , one also retrieves (2.21) in the particular case $J_0 = \emptyset$.

The associated transform algorithm. Like the direct method algorithm, the above methodology easily extends to an arbitrary number m of nontrivial constraints. The algorithm also consists of m one-dimensional integration steps. Possible pathological multiple poles are handled as in Section 2.4. Clearly, the computational complexity is described by n^m for the same reasons it is in the direct method.

The general case is better illustrated in Example 2.1 of Section 2.4. Setting $e_3 = (1, 1, 1)$, let $\Omega(ze_3) \subset \mathbb{R}^2$ be the polytope

$$\Omega(ze_3) := \{x \in \mathbb{R}_+^2 \mid x_1 + x_2 \leq z; -2x_1 + 2x_2 \leq z; 2x_1 - x_2 \leq z\},$$

whose area is $17z^2/48$.

Choose $\gamma_1 = \gamma_2 = 1$, $\gamma_3 = 2$, and $\lambda_3 = p - \lambda_2 - \lambda_1$, so that $\Re(p) = d = 4$ and

$$H(p) = \frac{1}{(2\pi i)^2} \int_{1-i\infty}^{1+i\infty} \int_{1-i\infty}^{1+i\infty} M(\lambda, p) d\lambda_1 d\lambda_2,$$

with

$$M(\lambda, p) = \frac{1}{\lambda_1 \lambda_2 (p - \lambda_1 - \lambda_2) (2p - \lambda_1 - 4\lambda_2) (2\lambda_1 + 3\lambda_2 - p)}.$$

We first integrate w.r.t. λ_1 . Only the real parts of the poles $\lambda_1 = 0$ and $\lambda_1 = (p - 3\lambda_2)/2$ are less than 1. Therefore, the residue of the 0-pole yields

$$\frac{1}{2\pi i} \int_{1-i\infty}^{1+i\infty} \frac{1}{\lambda_2 (p - \lambda_2) (2p - 4\lambda_2) (3\lambda_2 - p)} d\lambda_2, \quad (2.27)$$

whereas the residue of the $(p - 3\lambda_2)/2$ -pole yields

$$\frac{1}{2\pi i} \int_{1-i\infty}^{1+i\infty} \frac{4}{\lambda_2 (p - 3\lambda_2) (p + \lambda_2) (3p - 5\lambda_2)} d\lambda_2. \quad (2.28)$$

Applying again Cauchy residue theorem to (2.27) at the pole $\lambda_2 = 0$ (the only one whose real part is less than one) yields $-1/(2p^3)$.

Similarly, applying Cauchy residue theorem to (2.28) at the poles $\lambda_2 = 0$ and $\lambda_2 = -p$ (the only ones whose real part is less than one) yields $(4/3 - 1/8)/p^3$.

We finally obtain $H(p) = (4/3 - 1/8 - 1/2)/p^3 = 17/(24p^3)$, and so $h(z) = 17z^2/48$, the desired result.

Concerning the pathological case of poles on the integration path, we have already said that this issue occurs with probability zero in a set of problems with randomly generated data. Moreover, we have also proved that those poles are *removable singularities* and we only need to add all terms with poles on the integration path in order to get rid of them, with no need to implement the perturbation technique described in Section 2.4.

For example, had we chosen $\gamma_1 = \gamma_2 = \gamma_3 = 1$ (instead of $\gamma_1 = \gamma_2 = 1$ and $\gamma_3 = 2$) in the above example, we would have in both (2.27) and (2.28) a pole on the integration path $\{\Re(\lambda_2) = 1\}$, because $\Re(p) = 3$. But, adding (2.27) and (2.28) together yields

$$H(p) = \frac{1}{2\pi i} \int_{1-i\infty}^{1+i\infty} \frac{5p - 7\lambda_2}{\lambda_2 (p - \lambda_2) (2p - 4\lambda_2) (p + \lambda_2) (3p - 5\lambda_2)} d\lambda_2,$$

and, as expected, the integrand has no more poles on the integration path. In addition, only the real parts of poles $\lambda_2 = 0$ and $\lambda_2 = -p$ are less than 1, and their residues give $H(p) = (5/6 - 1/8)/p^3 = 17/(24p^3)$, and so $h(z) = 17z^2/48$.

The above two algorithms easily extend to the case $c \neq 0$ with slight ad hoc modifications, as one only replaces $A'\lambda$ with $A'\lambda - c$ in the denominator of the Laplace transform $\hat{F}(\cdot, c)$.

2.5 Notes

The computational complexity of volume computation is detailed in Dyer and Frieze [45] and in the nice exposition of Bollobás [24]. For more details on probabilistic approximate methods, the interested reader is referred to Bollobás [24], Dyer, Frieze, and Kannan [47], and Kannan, Lovász, and Simonovits [76]. For a general exposition of quasi-Monte Carlo methods for computing integrals (the volume being a particular one), see, e.g., Niederreiter [116]. Another method to approximate the volume of a convex polytope (with error estimates) is described in Kozlov [84]. It uses an appropriate integral formulation in which the integrand is written as a (truncated) Fourier series.

For an updated review of primal exact methods for volume computation and their computational complexity, see, e.g., Büeler et al. [29] and Gritzmann and Klee [62]. In particular, improved versions of some of the above algorithms are also described in [29], and the software package VINCI developed at ETH in Zürich offers several alternative methods working with the half-space description or the vertex description (or both) of Ω . It is available with free access at

<http://www.lix.polytechnique.fr/Labo/Andreas.Enge/Vinci.html>.

Most of the material in this chapter is from Barvinok [13], Büeler et al. [29], Brion and Vergne [27], and Lasserre and Zeron [95]. The residue method described in this chapter does not require one to work with a vector $c \in \mathbb{R}^n$ having some *regularity* property, as in Lawrence [102] and Brion and Vergne [27]. Despite this regularity property of c being *generic*, the resulting formula is not stable numerically as reported in [29].

Finally, let us mention that the Laplace transform approach developed in this chapter is also applicable to integration of other functions (e.g., Laurent polynomials) over other sets (e.g., ellipsoids). In some cases, one may even get the exact result in closed form or good approximations by truncated series; see, e.g., Lasserre and Zeron [96].



<http://www.springer.com/978-0-387-09413-7>

Linear and Integer Programming vs Linear Integration
and Counting

A Duality Viewpoint

Lasserre, J.B.

2009, XIV, 168 p. 2 illus., Hardcover

ISBN: 978-0-387-09413-7